Chapter 14

Real business cycles

Kurt Mitman

14.1 Introduction

This chapter introduces the concept of *business cycles* and presents the core framework for business cycle analysis: the real business cycle (RBC) model. This begs the question of what a business cycle is and why the qualification "real" is used for the model. Business cycles are the fluctuations of the economy around its long-run trend. The model is real, because we will be considering the stochastic version of the growth model developed in Chapter 7, without money or any nominal variables. This chapter can be seen as the natural continuation of Chapter 13, which focused on documenting and accounting for long-run growth patterns. Now, instead, we will be focused more on understanding short-run fluctuations in the economy.

The term business cycles does not necessarily imply that short-run fluctuations are periodic, like the sine or cosine functions. We begin the chapter by zooming in on the data to examine the short-run movements in aggregates like GDP. We'll discuss ways of extracting the trend from the longer time series and how to interpret the deviations from the trend in the short run. Then we'll document the co-movement of aggregate variables over the "cycle." Next, we'll show how the stochastic growth model can quantitatively be used to rationalize business-cycle fluctuations, the seminal contribution of Kydland and Prescott (1982). Finally, we'll discuss the successes (and shortcomings) of the RBC model, various extensions, and where the literature has gone since it was introduced.

14.2 Business cycles: a historical overview

The modern usage of the term business cycles refers to the fluctuations in economic activity that an economy experiences over a period of time. We refer to them as cycles because they are characterized by periods of economic expansion followed by periods of contraction or recession. The study of business cycles has its roots in studies of economic crises of the 19th century. The Panic of 1825 was arguably the first economic crisis unrelated to war or another external events, and resulted in a global economic downturn (Bordo (1998)). Interest in economic cycles grew as the 19th century saw a recurrence of boom and bust episodes throughout the industrialized world.

Early work was primarily descriptive in nature. Economists documented cyclical relationships in the data that appeared at different relative frequencies. Schumpeter (1939) summarizes the three main types of cycles that had been identified in the data (and named after the economists who identified them) during the late 19th and early 20th centuries:

1. Kitchin Inventory Cycle: A short cycle of about 3-5 years, often related to inventory adjustments.

2. Juglar Fixed Investment Cycle: A medium-term cycle of about 7-11 years, associated with fixed investment.

3. Kondratiev Long Waves: A long-term cycle lasting 45-60 years, driven by fundamental technological innovations.

While the early work did manage to identify cycles and provide evidence of co-movement of different economic variables, there was little distinction between the *impulses*, i.e., shocks that are at the origin of fluctuations, and the *propagation* of those shocks through the macroeconomy. One of the first to make a distinction between impulse and propagation was Swedish economist Knut Wicksell with his famous metaphor: "If you hit a rocking horse with a stick, the movement of the horse will be very different from the stick. The hits are the cause of the movement, but the system's own equilibrium laws condition the form of movement." The stick serves as the impulse to the economy, and the rocking in response to the hit the propagation.[1] In the 1930s, economists began to develop theories to explain both the impulses and their propagation throughout the economy.

Schumpeter's theory of business cycles was based on the idea that the impulses were caused by technological innovations developed by entrepreneurs. Innovations should be interpreted broadly here. For example, they include the invention of new products and methods of production. Since coming up with new ideas is an inherently random process, these served as the shocks to the economy. In his theory, Schumpeter argued that innovations tend to occur in clusters, with one breakthrough triggering a cascade of subsequent advances. This clustering of innovations gave rise to the cyclical nature of economic growth. Schumpeter viewed these cycles as fundamental to the capitalist process, with economic development driven by the entrepreneurial introduction of new technologies and methods. This idea of a fundamental link between growth and fluctuations continues to this day.[2] Along these lines, business cycle analysis today can to a large extent be thought of as (i) identifying the shocks that have occurred (and their origins) and (ii) how these shocks propagate through the economy.

[1] Frisch (1933) attributes the metaphor to Wicksell, but the quote cannot be found in a publication; it appears to originate from a conversation years earlier between Wicksell and another Swedish economist, Johan Åkerman.

[2] Schumpeter also emphasized the role of credit in facilitating innovations. Entrepreneurs may be constrained financially and have to borrow to finance their innovation activities; see Chapter 19. The banking system plays a crucial role in providing this credit. However, this can also lead to over-expansions of credit, contributing to the cyclical nature of the economy. These ideas are still relevant today: Ben Bernanke would later win the economics prize in memory of Alfred Nobel for his work linking the severity of the Great Depression to bank failures and credit disruptions.

Schumpeter's theory was, in some sense, "creatively destroyed" before it made a significant impact on economic thought at the time. The Great Depression of the 1930s prompted a rethinking of economic theories that could explain such severe economic downturns. John Maynard Keynes was at the forefront of this intellectual revolution. Keynes argued that insufficient aggregate demand could lead to prolonged periods of high unemployment. He believed that in the face of a decrease in aggregate demand, wages and prices might not adjust quickly enough to restore full employment. As a result, government intervention, in the form of fiscal policy, was necessary to stabilize the economy. One reason, perhaps, why Keynes's theory gained traction was that it both explained the Great Depression and prescribed a way to fight the global slump. His seminal work, "The General Theory of Employment, Interest, and Money" (Keynes (1936)), laid the foundation for what came to be known as Keynesian economics. This school of thought dominated the study of business cycles until the 1970s.

In the late 1960s and 1970s, the economic landscape changed with the occurrence of stagflation, a combination of high inflation and high unemployment. This phenomenon was difficult to explain using traditional Keynesian theories that posited a negative relationship between unemployment and inflation. Further, the surge in oil prices following OPEC restrictions on oil supply pointed toward supply-side explanations, which were largely viewed as second-order in the Keynesian theory. At the same time, monetary and fiscal policies in advanced economies during the 1970s appeared ineffective in achieving key policy targets of low and stable inflation. As the Keynesian paradigm struggled to explain stagflation, some economists began to challenge its methodological foundations. Milton Friedman, Edmund Phelps, and Robert Lucas argued that the relationship between inflation and economic activity would change as inflation expectations respond to changes in macroeconomic policy. This type of argument against using reduced-form relationships came to be known as the Lucas Critique (Lucas (1976)).

In a series of influential papers, Lucas also advocated for an approach based on rational expectations: the assumption that individuals understand the structure of the economy and revise their expectations accurately in response to changes in policy. In order to implement such an approach, attention turned to models with microfoundations on the grounds that specifying the primitives of the model in terms of policy-invariant features such as preferences and technologies would avoid the Lucas Critique.

Lucas and Rapping (1969) took a first step towards incorporating rational expectations into business cycle analysis and made significant contributions to the study of labor supply in the context of business cycle fluctuations. Central to Lucas and Rapping's theory is the idea of intertemporal substitution. They posited that individuals make labor supply decisions based on expected future wages. When individuals expect higher wages in the future, they are willing to substitute leisure today for work in the future, and vice versa. In their model, temporary, unanticipated changes in wages lead workers to adjust their labor supply. They argued that observed fluctuations in employment could be largely explained by workers' voluntary decisions to adjust their labor supply in response to unexpected wage changes, rather than by involuntary unemployment. This perspective was in contrast to the traditional Keynesian view, which emphasized the role of demand deficiencies in causing unemployment during recessions. Lucas and Rapping's work suggested that labor market fluctuations might be better understood by focusing on the economy's supply side and workers' intertemporal decisions. Intertemporal substitution has remained central in modern macroeconomic models.

While Lucas emphasized the need for micro-founded macro models, it was Kydland and Prescott who operationalized this vision as a full model of the economy with the development of real business cycle (RBC) theory. Building on rational expectations and intertemporal substitution, they used a stochastic version of the neoclassical growth model, arguing that business cycle fluctuations result primarily from real shocks—especially productivity shocks—rather than policy interventions. Their work emphasized the importance of understanding the intertemporal decisions of households and firms in analyzing macroeconomic phenomena. From a methodological point of view, Kydland and Prescott also introduced into economics the concept of solving models numerically and relying on computer simulations to analyze model behavior (Kydland and Prescott, 2004). By calibrating and simulating dynamic general equilibrium models, they demonstrated that such shocks could replicate observed business cycle patterns. Their approach marked a methodological shift: numerical solution and simulation became central tools of macroeconomic analysis. Prescott viewed the RBC model as a foundational framework for macroeconomics, akin to supply and demand in microeconomics. This work established the field of quantitative macroeconomics and became the dominant paradigm.

14.3 A first look at the data

One of the key methodological innovations of the RBC paradigm was to seriously connect a micro-founded business cycle model with real-world data. To begin, we examine the time series of quarterly real U.S. GDP in Figure 14.1, plotted for the post-war period. As emphasized in Chapter 13, U.S. GDP displays a strong upward growth trend. However, in this chapter we will focus on fluctuations around that trend, or the business cycles. Two natural questions arise: how do we define business cycles in the data, and how do we isolate the cycle from the trend component in GDP and other macroeconomic variables?

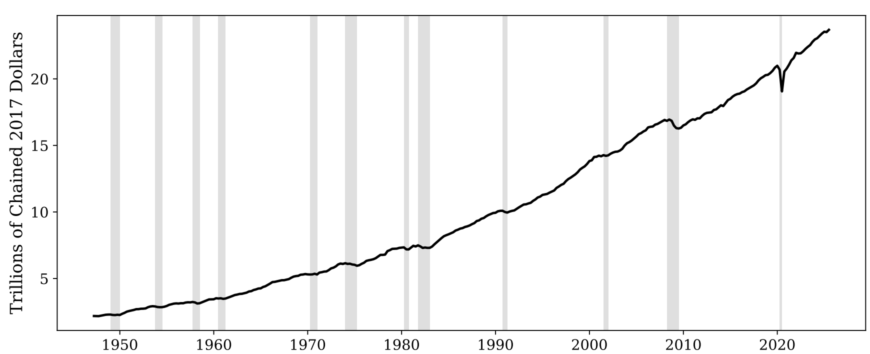

Figure 14.1: U.S. Real GDP and NBER Recessions

In the U.S., the business cycle is defined by its turning points, namely the *peaks* and *troughs*. In between a peak and a trough is a *recession*, and between a trough and peak an *expansion*. Periods of expansion are considered "normal times," whereas recessions are thought of as brief episodes of economic contraction. While the terms seem clear, they

The Business Cycle

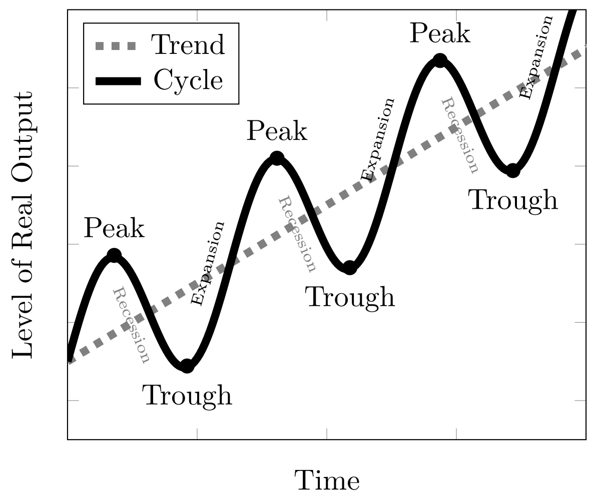

Figure 14.2: Stylized representation of business cycles

are not formal statistical definitions that can be applied to real-time data to determine the current state of the economy. Because of the underlying trend of economic growth, a recession could still be a period of positive GDP growth, albeit at a slower rate. This example highlights the importance of separating the cycle from the trend in the economy. In the U.S., a committee officially determines whether the economy is in a recession. The NBER Business Cycle Dating Committee has officially dated U.S., recessions since 1978. The committee uses a variety of monthly and quarterly data on measures of economic activity to date business cycles retrospectively. The description of how they determine when the economy is in recession is reproduced below:[3]

NBER Recessions

The NBER's definition emphasizes that a recession involves a significant decline in economic activity that is spread across the economy and lasts more than a few months. In our interpretation of this definition, we treat the three criteria—depth, diffusion, and duration—as somewhat interchangeable. That is, while each criterion needs to be met individually to some degree, extreme conditions revealed by one criterion may partially offset weaker indications from another. For example, in the case of the February 2020 peak in economic activity, the committee concluded that the subsequent drop in activity had been so great and so widely diffused throughout the economy that, even if it proved to be quite brief, the downturn should be classified as a recession. Because a recession must influence the economy broadly and not be confined to one sector, the committee emphasizes economy-wide measures of economic activity.

Given the inherently qualitative nature of the NBER recession dating, we will instead focus on a statistical definition following Lucas (1977) and Kydland and Prescott (1982)

[3]Source: https://www.nber.org/research/business-cycle-dating, accessed January 5, 2024.

that defines business cycles as recurrent fluctuations of output along a slow-moving trend, and the associated co-movements of other aggregate quantities. The basic idea is that we want to split trending or non-stationary variables, such as GDP, into a non-stationary trend component and a stationary, cyclical component. A stylized representation of the decomposition into trend and stationary is plotted in Figure 14.2. Through the lens of the neoclassical growth model presented in the previous chapter, the goal would be to remove the exogenous deterministic labor-augmenting productivity growth, and transform the model into a stationary one. In the context of the model, that would be easy to achieve. However, identifying the trend in the data presents more of a challenge, since we only observe the combination of trend and stationary components. Further complicating the analysis is the fact that the longer-run trend may not be constant, but could be something more akin to the Kondatiev Long Wave mentioned above. To make progress we will discuss techniques for *filtering* the data, that is, extracting trends of different frequencies. In this chapter, we will discuss the most common method for detrending in the RBC literature, the Hodrick-Prescott filter, as well as another commonly used method, the band-pass filter.

14.3.1 Filters

The most common approach to business cycle analysis involves filtering the data to transform it into a stationary series that can be studied in isolation. The idea is to come up with a graph similar to Figure 14.3.

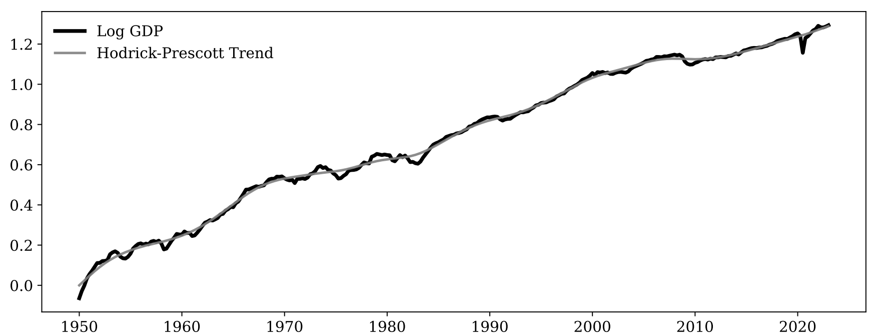

Figure 14.3: Log of U.S. GDP: data and trend from Hodrick-Prescott filter

The optimal way of extracting the trend and cycle components from the data should, in principle, depend on the underlying theory or assumptions about the data generating process (DGP). We can represent the DGP for a macroeconomic time series, Y_t, as a trend-stationary process by assuming that we can decompose them into trend and cycle components, $Y_t = \tau_t + Y_{c,t}$. For example, imagine the trend component is simple exogenous growth $\tau_t = y_0 e^{\gamma t}$ and the cyclical component is simply i.i.d, $Y_{c,t} = \epsilon_t$. One could then simply estimate a linear trend in $\log Y_t$ to capture the trend component, and the residual from that trend would be represent the cyclical one. Another common formulation for the DGP is to assume that the process is difference-stationary. A process that's stationary in first-differences could be represented as $\Delta Y_t = \bar{Y} + \phi(L)\epsilon_t$, where L is the lag operator ($LX_t = X_{t-1}$), $\Delta \equiv (1 - L)$ is the difference operator, and ϕ is a polynomial. For example, a random walk with constant drift could be represented as $\Delta Y_t = \bar{Y} + \epsilon_t$. In finite time series the data can be approximated equally well by difference-stationary and trend-stationary processes (Hamilton, 1994). In practice, therefore, there are many potential DGPs that could be consistent with the time series data. Any filter chosen will be understood to only represent an approximation of the true underlying DGP, and thus will only capture some aspects of the data. Most business cycle analysis therefore does not specify an underlying DGP for the data, but instead proceeds with filtering in a more informal way. Proceeding in this fashion is appropriate if the results are not sensitive to the particular filter chosen. While this type of robustness analysis was carried out in early RBC contributions (e.g., Prescott, 1986), it is typically omitted in current business cycle analysis. This chapter focuses primarily on the most widely-used filter in business cycle analysis, the Hodrick-Prescott filter.[4]

The Hodrick-Prescott Filter Given a time series $\{Y_t\}_{t=0}^{T}$, the Hodrick-Prescott (HP) filter decomposes it into a trend component τ_t and a cyclical component $Y_{c,t}$ such that

$$Y_t = \tau_t + Y_{c,t}.$$

The trend component τ_t is obtained by minimizing the objective function

$$\min_{\{\tau_t\}} \sum_{t=1}^{T} (Y_t - \tau_t)^2 \tag{14.1}$$

subject to

$$\sum_{t=2}^{T-1} \left[(\tau_{t+1} - \tau_t) - (\tau_t - \tau_{t-1})\right]^2 \leq \mu,$$

where μ is a parameter that controls the smoothness of the trend component. Smaller values of μ impose smoother paths for the trend component of the time series. In the limit, with $\mu = 0$, $\tau_{t+1} - \tau_t = \tau_t - \tau_{t-1}$ $\forall t$, which imposes a linear trend. Thus, the minimization balances a trade-off between tracking the raw data perfectly and following a linear trend.

Under the assumption that the constraint binds, we can re-write the minimization problem using the Lagrangian, where we attached λ as the Lagrange multiplier on the constraint:

$$\min_{\{\tau_t\}} \left[\sum_{t=1}^{T} (Y_t - \tau_t)^2 + \lambda \sum_{t=2}^{T-1} \left((\tau_{t+1} - \tau_t) - (\tau_t - \tau_{t-1})\right)^2\right]$$

.

In this formulation, μ no longer appears explicitly, and the filter can instead be parameterized by λ, which serves as a smoothing parameter. Higher values of λ impose tighter smoothness constraints on the trend component. In the limit as $\lambda \to \infty$, the extracted trend converges to a linear function. For quarterly data, Prescott (1986) suggested a value of $\lambda = 1{,}600$ as a reasonable choice. However, the appropriate value of λ depends on the frequency of the data. Ravn and Uhlig (2002) show that the HP filter parameter should be scaled by the fourth power of the ratio of observation frequencies. For instance, if

[4]See Hamilton (2018) for a discussion of shortcomings of the Hodrick-Prescott filter.

λ = 1,600 is used for quarterly data, the corresponding value for annual data would be $1{,}600 \times (1/4)^4 = 6.25$.

To implement the HP filter, one needs to solve the above optimization problem. While there are various methods to do this, a common approach is to use matrix algebra and the system of first-order conditions (FOCs) from the problem. The FOC for τ_t for $2 \leq t \leq T-2$ are given by

$$\tau_t - \lambda\left(\Delta\tau_t - \Delta\tau_{t-1} - 2\left(\Delta\tau_{t+1} - \Delta\tau_t\right) + \Delta\tau_{t+2} - \Delta\tau_{t+1}\right) = Y_t.$$

For the end conditions the FOCs will be slightly different (because they cannot rely on data from before $t = 0$ or after $t = T$). But one can show that this system of FOCs has the following form: $(I + \lambda A)\boldsymbol{\tau} = \boldsymbol{Y}$, where I is the identity matrix, $\boldsymbol{\tau}$ is the vector of τ_t, $\boldsymbol{Y}$ is the vector of Y_t, and A is a square pentadiagonal matrix given by

$$\begin{bmatrix}
1 & -2 & 1 & 0 & \cdots & \cdots & \cdots & \cdots & 0 \\
-2 & 5 & -4 & 1 & 0 & \cdots & \cdots & \cdots & 0 \\
1 & -4 & 6 & -4 & 1 & 0 & \cdots & \cdots & 0 \\
0 & 1 & -4 & 6 & -4 & 1 & 0 & \cdots & 0 \\
\vdots & & \ddots & \ddots & \ddots & \ddots & \ddots & \vdots & \\
0 & 0 & \cdots & 1 & -4 & 6 & -4 & 1 & 0 \\
0 & 0 & \cdots & 0 & 1 & -4 & 6 & -4 & 1 \\
0 & 0 & \cdots & 0 & 0 & 1 & -4 & 5 & -2 \\
0 & 0 & \cdots & 0 & 0 & 0 & 1 & -2 & 1
\end{bmatrix}$$

Given a value of λ, we can obtain the trend by solving the system $\boldsymbol{\tau} = (I + \lambda A)^{-1}\boldsymbol{Y}$. Then, we can recover the cyclical component by differencing the data from the trend component. In practice, most common programming languages used by economists (e.g., MATLAB and Python) have built-in functions for HP-filtering the data. There are some things to keep in mind, however, when using the HP filter. First, the choice of λ can influence the resulting decomposition. There is no universally optimal choice for λ, and different applications may require different values (see Marcet and Ravn, 2004, for further discussion). Second, the filter may produce spurious results towards the beginning and end of the sample. Finally, the particular functional form for the HP-filter is not based on economic theory.

Spectrum and band-pass filters In our discussion so far of filtering and separating trends versus cyclical components, we have discussed the idea of trying to isolate cycles at "business cycle frequencies," that is, cycles that occur with a frequency of something like 5-10 years. Ideally, then, we'd like to extract the components of the data at that frequency and remove any longer-run trends that occur with longer frequencies. We typically express time-series data in economics in the time domain, but for stationary time series there exists an equivalent representation in the *frequency domain*, as the sum of sine and cosine functions of various frequencies, amplitudes and phases.[5] Higher frequency means a rapidly changing

[5]To remind you of the properties of sine, the function $\sin(2\pi t f)$ is plotted for difference frequencies f in Appendix Figure 14.A.1. The first sine wave completes a full cycle at $t = 2\pi$, whereas the wave with frequency of 2 completes two full cycles by that time, and the wave of frequency 4 completes four cycles.

wave. The sine function has a range of $(-1, 1)$, but the *amplitude* can be adjusted by pre-multiplying the function. The *phase*, or where the sine wave is in its cycle, can be adjusted by adding or subtracting from the argument of the function (which shifts the wave to the left or right). Thus, we can write a general sine wave as $B\sin(ft+\phi)$, where B controls the amplitude, f controls the frequency, and ϕ controls the phase.

It is possible to represent any stationary process Y_t as follows:

$$Y_t = \int_0^\pi A(\omega)\cos(\omega t)d\omega + \int_0^\pi B(\omega)\sin(\omega t)d\omega,$$

where for any ω, $A(\omega)$ and $B(\omega)$ are random variables, and those functions A and B are what define the process. What this expression implies is that the time series is a sum, with random weights, of sine and cosine waves, and the sum is over all frequencies between 0 and π. If we want to isolate a particular range of frequencies, we can imagine a filter that removes any frequency that falls outside of the intended range. Such a filter exists and is called the *band-pass* filter. Using the band pass filter one can isolate and extract cyclical components focusing on any specific range of frequencies (e.g., 5-10 years for business cycles) in the data. A more detailed discussion of the band-pass filter is provided in Appendix 14.A.

14.3.2 Stylized business-cycle facts

Having reviewed tools for extracting business-cycle components, we now present key empirical regularities that motivated the RBC framework. These stylized facts are based on filtered data, primarily using the HP filter, but we also demonstrate their robustness to alternative methods such as first differencing and the band-pass filter. In this section we'll also investigate how the business cycle facts have changed over time, comparing the original statistics from Kydland and Prescott (1982) to those based on all currently available data. We begin with unconditional moments—specifically, second moments and correlations of aggregate time series. Later, we will examine conditional correlations, following TFP shocks.

One of the defining features of business cycles is the *comovement* of aggregates in booms and recessions. The standard practice in the business-cycle literature is to define co-movements by looking at correlations of aggregates relative to GDP, $Y_{c,t}$. Focusing on contemporaneous correlations, we say that a variable $X_{c,t}$ is:

- *Pro-cyclical* if it comoves positively with output, $corr(Y_{c,t}, X_{c,t}) > 0$.
- *Countercyclical* if it comoves negatively with output, $corr(Y_{c,t}, X_{c,t}) < 0$.

In addition to contemporaneous correlations, to understand propagation it is also helpful to understand which variables *lead* or *lag* movements in output. A variable $X_{c,t}$ is said to be leading output if $corr(Y_{c,t+s}, X_{c,t})$ is highest and positive for some $s > 0$, and is lagging if $corr(Y_{c,t+s}, X_{c,t})$ is highest and positive for some $s < 0$.

Data sources

The data on output, consumption, investment, government consumption, and the price level (here measured as the GDP deflator) come from the National Income and Product Accounts

(NIPA) provided by the U.S. Bureau of Economic Analysis (BEA). The data on employment come from the Current Employment Statistics provided by the U.S. Bureau of Labor Statistics (BLS). The data on the unemployment rate is from the Current Population Survey (CPS) provided by the BLS. The data on hours and wages (per hour) are from the non-farm business sector provided by the BLS. The fed funds rate (the policy rate for the central bank) is provided by the Board of Governors of the Federal Reserve System. All quantity variables are expressed in per capita terms by dividing by the civilian non-institutional population aged 16+. All variables are all expressed in logs, except for the federal funds rate. The data are at a quarterly frequency, and all cyclical components are then extracted using an HP filter with $\lambda = 1{,}600$. We plot the data and HP trend for output, consumption, investment and, hours in Figure 14.4.

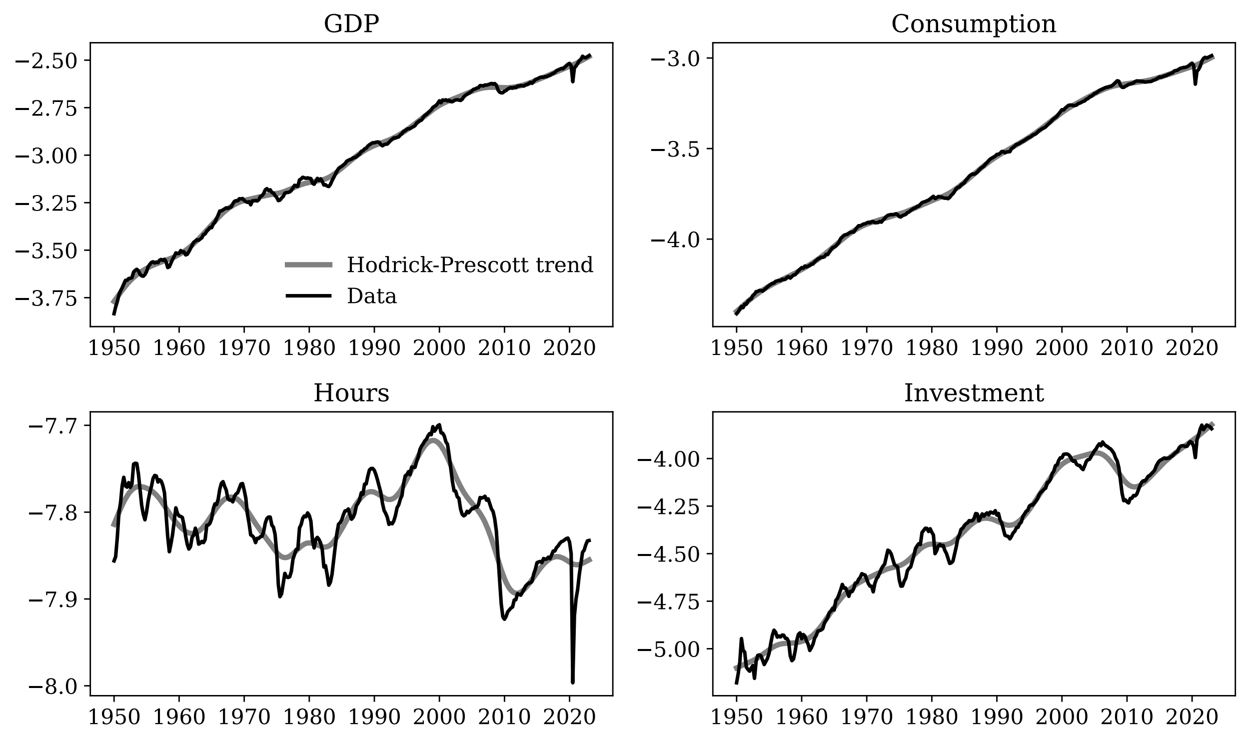

Figure 14.4: Actual and trends of logs of U.S. aggregates

Business cycle facts

The second moments and correlations for the U.S. are listed in Table 14.1 and the main stylized facts are summarized below.[6]

Fact 1 **Consumption smoothing:** Consumption is less volatile than output, with a relative standard deviation of 0.65.

[6] In Chapter 23, we show that many of these stylized facts hold across countries.

Table 14.1: Business cycle moments U.S. Data 1949-2022

Variable x	Standard Deviation (%)	Relative Std to σ_y	Auto-correlation	Cross correlation of x with $y(t-1)$	$y(t)$	$y(t+1)$
Output (y)	1.63	1.00	0.78	0.78	1.00	0.78
Consumption	1.07	0.65	0.64	0.58	0.75	0.53
Gov. Consumption	3.02	1.85	0.89	0.23	0.15	0.04
Investment	4.37	2.68	0.86	0.61	0.77	0.71
Employment	1.58	0.97	0.81	0.79	0.79	0.48
Hours	2.11	1.29	0.80	0.77	0.86	0.61
Unemployment	15.63	9.57	0.80	-0.79	-0.83	-0.55
Lab. Productivity	1.15	0.70	0.72	-0.03	0.27	0.38
Wages	1.24	0.76	0.72	-0.03	-0.00	0.15
Price Level	0.92	0.56	0.92	-0.02	-0.10	-0.21
TFP	0.90	0.55	0.75	0.17	0.51	0.52
Fed Funds Rate	3.60	2.20	0.97	0.23	0.18	0.08

Fact 2 **High investment volatility:** Investment is the most volatile component of GDP, with a standard deviation about 2.7 times that of output.

Fact 3 **Labor market dynamics:** Hours worked have similar volatility to output; employment is slightly less volatile, while unemployment is highly volatile and strongly countercyclical.

Fact 4 **Cyclicality:** Consumption, investment, employment, hours, and productivity are all procyclical, while unemployment is strongly countercyclical and lags output.

Fact 5 **Price and wage dynamics:** Despite large swings in quantities, real wages and the price level are much less volatile and less cyclical.[7]

Fact 6 **High serial correlation:** All major macroeconomic aggregates exhibit significant serial correlation. Notably, TFP and labor productivity are about as persistent as output itself.

In addition, if one disaggregates production by sector, it is possible to show that most sectors move together over the cycle, though some (e.g., mining) may behave differently. Disaggregating consumption reveals that durable goods are most volatile, followed by non-durables, with services being the least volatile.[8]

[7]One of the difficulties in measuring the cyclicality of wages is due to selection. The workers who lose their jobs in recessions are not randomly selected. On average, those workers tend to be lower wage workers. Thus, in recessions the composition of workers shifts towards higher wage workers, mitigating the decline in the average wage.

[8]To show that the stylized business cycle facts are robust to different filtering, in Appendix Tables 14.A.1 and 14.A.2 the analysis is repeated using a band-pass filter and the first difference filter. While the exact numbers clearly are not identical across the different filtering methods, the volatilities and relative volatilities are similar. And, importantly, the cyclicality of all variables (except for wages) is the same across all three filters.

The high degree of serial correlation observed in all major macroeconomic aggregates raises an important question: Does this persistence reflect highly persistent shocks to the economy, or do even transitory shocks produce long-lasting effects because of strong internal propagation mechanisms? The fact that TFP and labor productivity are about as persistent as output suggests that models driven primarily by TFP shocks will naturally generate persistent fluctuations in other aggregates, even in the absence of substantial endogenous propagation.

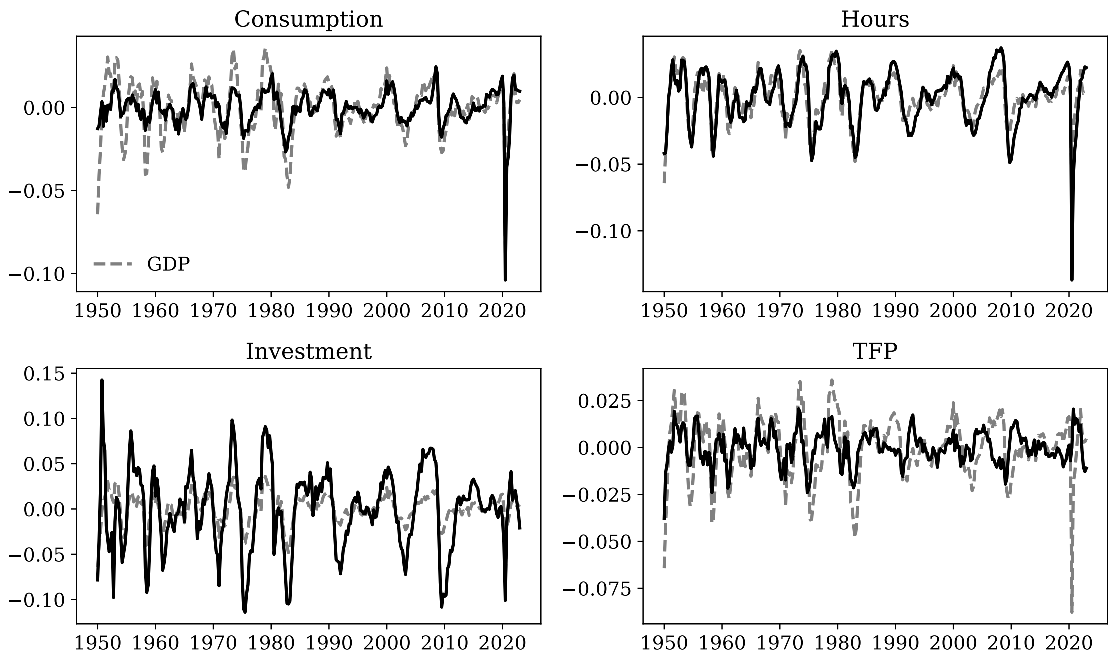

Figure 14.5: Deviations from trend of U.S. aggregates

To illustrate business cycle co-movement, Figure 14.5 plots consumption, investment, hours, and TFP as deviations from trend. In each panel, GDP is shown as a dashed line to highlight both co-movement and relative volatility across series. Notably, the lower right panel shows that, after the mid-1980s, TFP and GDP appear less correlated. To investigate the potential changing nature of business cycles, in Tables 14.2 and 14.3 we recompute the cyclicality using two subsamples, one from 1949-1984 (corresponding roughly to the time period of the original RBC papers), and one for the 1985-2022.

Output is significantly less volatile after the 1980s as compared to before (standard deviation of 1.24% compared to 1.97%); this is sometimes referred to as the "Great Moderation." While the strong co-movement of output, consumption, investment, and labor market variables has been stable over time, some aggregates exhibit significantly different co-movement since the original RBC papers were written. In particular, government consumption switches from being pro-cyclical in the early period to being countercyclical in the second half of the sample. This suggests that fiscal policy has perhaps become more Keynesian over time, using government spending to try to stabilize recessions. The federal funds rate has also become

Table 14.2: Business cycle moments U.S. data 1985-2022

Variable x	Standard Deviation (%)	Relative Std to σ_y	Auto-correlation	Cross correlation of x with $y(t-1)$	$y(t)$	$y(t+1)$
Output (y)	1.24	1.00	0.64	0.64	1.00	0.64
Consumption	1.20	0.97	0.56	0.56	0.86	0.43
Gov. Consumption	1.26	1.02	0.83	-0.32	-0.38	-0.39
Investment	3.64	2.94	0.90	0.63	0.79	0.70
Employment	1.56	1.26	0.72	0.69	0.78	0.34
Hours	2.19	1.77	0.74	0.70	0.88	0.52
Unemployment	15.60	12.59	0.74	-0.72	-0.83	-0.43
Lab. Productivity	1.09	0.88	0.76	-0.47	-0.29	-0.05
Wages	1.46	1.18	0.69	-0.21	-0.31	-0.06
Price Level	0.68	0.55	0.93	0.37	0.30	0.14
TFP	0.73	0.59	0.78	-0.33	-0.03	0.10
Fed Funds Rate	2.83	2.29	0.99	0.30	0.29	0.25

more pro-cyclical, suggesting a change in monetary policy over the time period. Relatedly, the price level goes from being countercyclical to pro-cyclical. Perhaps more surprisingly, wages and labor productivity switch from being strongly pro-cyclical in the early sample to counter-cyclical in the later sample. Indeed, the recessions from 1991 onwards featured so-called "jobless recoveries" in which output recovered quickly, but employment was much slower to recover. Jobless recoveries and countercyclical labor productivity present a challenge for models of productivity-driven business cycles that we will discuss towards the end of the chapter.

14.3.3 Conditional data moments

As discussed at the beginning of the chapter, modern business cycle analysis is interested both in understanding the impulses to the economy and their propagation. A fundamental tool for understanding the propagation of structural shocks is through an *impulse response function* or IRF introduced in Chapter 3.5.2. The IRF traces out the effect of a one-time shock to the economy to current and future values of other aggregates, i.e., the conditional response of variables to the realization of the shock. IRFs, which are linear representations and should be interpreted as approximations applicable to small shocks, visualize the complex interdependencies and dynamic responses in the macroeconomy. Analyzing the IRF is a key tool for understanding the dynamics in the system, and as we'll discuss in the next section we can construct model analogues to IRFs to directly compare model and data. Below, we apply the two most common methods, discussed in detail in Chapter 8, for estimating conditional data moments and IRFs to the RBC context: vector autoregressions (VARs) and local projections (LPs).

Vector autoregression A VAR model for business cycles might be used to forecast labor productivity, output, and investment simultaneously, where each of these variables is expected to affect the others:

Table 14.3: Business cycle moments U.S. data 1949-1984

Variable x	Standard Deviation (%)	Relative Std to σ_y	Auto-correlation	Cross correlation of x with $y(t-1)$	$y(t)$	$y(t+1)$
Output (y)	1.97	1.00	0.84	0.84	1.00	0.84
Consumption	0.89	0.45	0.81	0.66	0.75	0.69
Gov. Consumption	4.15	2.11	0.90	0.35	0.26	0.12
Investment	5.00	2.54	0.83	0.59	0.76	0.70
Employment	1.56	0.79	0.90	0.89	0.83	0.59
Hours	2.00	1.02	0.88	0.88	0.90	0.70
Unemployment	15.66	7.96	0.87	-0.89	-0.88	-0.66
Lab. Productivity	1.21	0.61	0.68	0.23	0.63	0.66
Wages	0.93	0.47	0.77	0.13	0.33	0.40
Price Level	1.14	0.58	0.91	-0.17	-0.28	-0.37
TFP	1.05	0.54	0.73	0.42	0.76	0.73
Fed Funds Rate	3.89	1.98	0.95	0.25	0.16	0.01

$$
\begin{aligned}
A_t &= a_1 + b_{11}A_{t-1} + b_{12}y_{t-1} + b_{13}i_{t-1} + \eta_{1,t} \\
y_t &= a_2 + b_{21}A_{t-1} + b_{22}y_{t-1} + b_{23}i_{t-1} + \eta_{2,t} \\
i_t &= a_3 + b_{31}A_{t-1} + b_{32}y_{t-1} + b_{33}i_{t-1} + \eta_{3,t}.
\end{aligned}
$$

Defining

$$X_t \equiv \begin{bmatrix} A_t \\ y_t \\ i_t \end{bmatrix}, \boldsymbol{\epsilon}_t \equiv \begin{bmatrix} \eta_{1,t} \\ \eta_{2,t} \\ \eta_{3,t} \end{bmatrix},$$

we can write the (reduced-form) VAR as:

$$X_t = \boldsymbol{B}X_{t-1} + \boldsymbol{\eta}_t,$$

where $\boldsymbol{B}$ is the matrix of coefficients on the lags (in many empirical implementations more than one lag is used) and $\boldsymbol{\eta}_t$ is the vector of innovations. As discussed in Chapter 8, the key identification issue is what we will assume about the relationships among contemporaneous ηs. To implement theoretical restrictions, we first write the (structural) VAR as

$$\boldsymbol{A}X_t = \boldsymbol{F}X_{t-1} + \boldsymbol{\epsilon}_t,$$

where $\boldsymbol{A}$ now is a matrix that captures the contemporaneous relationships among the variables in X_t and ϵ_t is a vector of structural shocks (as opposed to just reduced-form innovations). Let us, as also proposed as an example in Chapter 8, use timing restrictions:

$$\boldsymbol{A} = \begin{bmatrix} a_{11} & 0 & 0 \\ a_{21} & a_{22} & 0 \\ a_{31} & a_{32} & a_{33} \end{bmatrix}.$$

Thus, we assume that the first variable (TFP) does not respond to contemporaneous movements in output (since the off diagonal elements in the first row are zero). The second variable, output, responds to productivity, but not to investment. Finally, investment can respond contemporaneously to both productivity and output. To solve the model, we invert the A matrix, then estimate the reduced-form VAR, and then back out the structural shocks from the reduced-form errors.[9] The impulse response of variable i at time $t+s$ to a shock to variable j at time t can then be backed out by solving for $\partial X_{i,t+s}/\partial \epsilon_{j,y}$.

The identifying restrictions used here align with the RBC model presented in the next section: if TFP shocks drive business cycles and capital is predetermined, it is reasonable to assume that productivity does not immediately respond to output or investment. Likewise, since capital investment takes a period to become productive, output and productivity should not respond to current investment. Figure 14.6 shows the impulse responses to a TFP shock based on an SVAR estimated using U.S. data. The IRFs show significant co-movement between productivity, output, and investment from the shock to TFP. These IRFs suggest that shocks to productivity are promising candidates for explaining the unconditional moments described in the preceding section. Productivity shocks could generate the right co-movements and also the relative volatilities between output and investment.

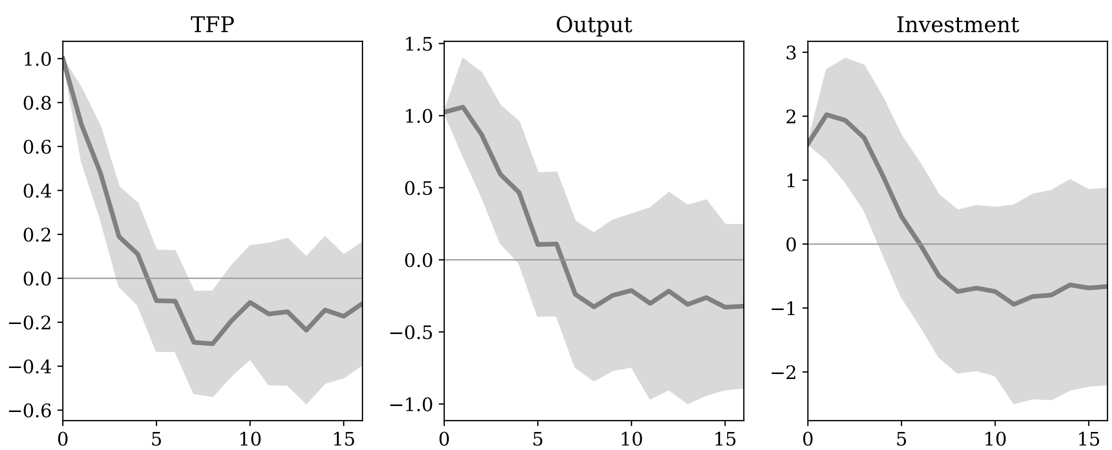

Figure 14.6: Impulse responses to TFP shock estimated with structural VAR

Local projections In the VAR above, we used timing restrictions to identify TFP shocks. If we assume that our measure of TFP (calculated as the Solow residual) is an exogenous shock, we can also estimate IRFs to these shocks by means of local projections. The results are plotted in Figure 14.7. The impulse responses are qualitatively similar to those derived from the SVAR, which is reassuring and further reinforces the notion that TFP shocks are a potential important driver of business cycle fluctuations. We now turn to the theory behind the real business cycle model.

[9] See the box on recursive identification in Chapter 8.

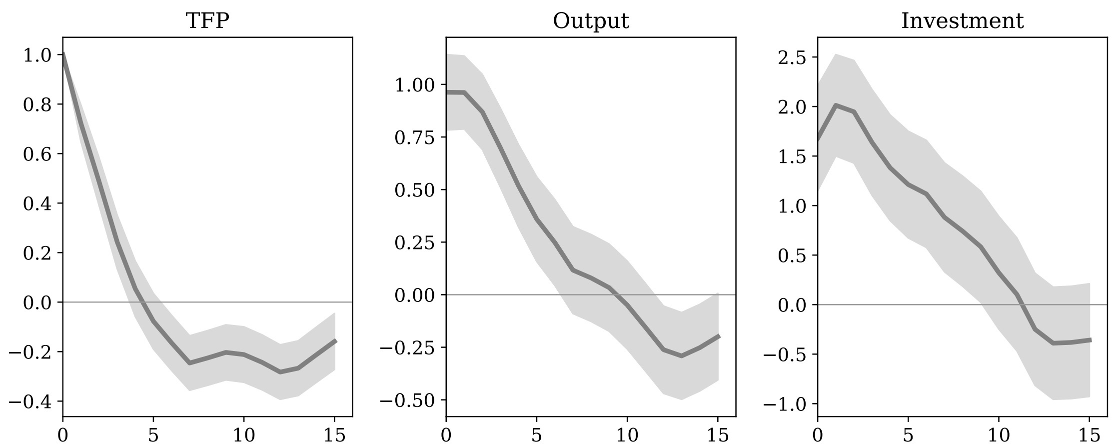

Figure 14.7: Impulse responses to TFP shock estimated with local projection

14.4 Real business cycle models

To develop a micro-founded framework for business cycle analysis, Kydland and Prescott extended the neoclassical growth model from Chapter 4 by introducing stochastic TFP, giving rise to the real business cycle (RBC) model outlined in Chapter 7. This stochastic neoclassical growth model is the foundation of modern business cycle theory and serves as the core for both New Keynesian (Chapter 18) and International Macro models (Chapter 23). Before presenting the full model with capital, we begin with a simplified version that includes only labor supply. This version admits a closed-form solution and serves to illustrate why labor alone cannot account for the key business cycle facts.

14.4.1 A simple business cycle model

Consider a representative household that maximizes its expected discounted utility. Time is discrete and lasts forever. The utility function is characterized by log preferences that are separable between consumption c_t and leisure $1-\ell_t$, where ℓ_t represents labor supply. The utility function is given by

$$U = E_0 \sum_{t=0}^{\infty} \beta^t \Big(\ln(c_t) + \phi \ln(1-\ell_t) \Big). \tag{14.2}$$

The production function of the representative firm is linear in labor and subject to TFP shocks z_t:

$$Y_t = z_t \ell_t, \tag{14.3}$$

where we assume that z_t follows an AR(1) process in logs, $\log z_t = \rho \log z_{t-1} + \sigma\epsilon$. Perfect competition and linear production on the firm side imply that the equilibrium wage will be equal to TFP, $w_t = z_t$.

The representative household receives income from supplying labor and has access to trade a risk-free one period bond that is in zero net supply:

$$c_t + B_t = w_t \ell_t + R_t B_{t-1}, \tag{14.4}$$

where R_t is the gross real rate and B_t are the bonds purchased by the household in period t.

Equilibrium and labor supply In equilibrium, the household's choice of labor supply balances the marginal utility of consumption against the disutility of labor. Substituting out for consumption using the budget constraint, and taking the first-order condition with respect to ℓ_t, we arrive at the static intratemporal optimization condition

$$\frac{w_t}{c_t} = \frac{\phi}{1 - \ell_t}. \tag{14.5}$$

We can derive the Euler equation from the intertemporal optimality condition:

$$\frac{1}{c_t} = \beta \mathbb{E}_t \frac{R_{t+1}}{c_{t+1}}. \tag{14.6}$$

Constant labor supply is optimal in equilibrium The last step in characterizing the equilibrium is to recognize that since the bond is in zero net supply and we have a representative household, it must be the case that in equilibrium the household holds no bonds. Next, we can eliminate consumption from the intratemporal labor supply condition using the budget constraint and setting $B_t = 0$ for all t to arrive at

$$\frac{\phi}{1 - \ell_t} = \frac{w_t}{w_t \ell_t} = \frac{1}{\ell_t}. \tag{14.7}$$

In this setup, the equilibrium labor supply is constant across periods, independent of the wage, and determined solely by preferences. Although the economy experiences fluctuations, there is no amplification or propagation, only immediate responses to shocks. Since output equals fixed labor input times TFP, output, consumption, and TFP are perfectly correlated, while hours worked remain unchanged. This is clearly inconsistent with the business cycle facts presented earlier in the chapter. At first glance, this result may seem puzzling. As shown in Chapter 12, log-log preferences imply a Frisch elasticity of $e_F = (1-\ell)/\ell$. If households work one-third of their available time, this gives a Frisch elasticity of 2—suggesting highly elastic labor supply. So why is there no response of hours to fluctuations in the real wage? The key lies in the definition of the Frisch elasticity: it measures labor supply responses holding the marginal utility of consumption constant. In this model, households cannot transfer resources across time, as there are no savings or borrowing instruments. Thus, the real interest rate must adjust to make it optimal for households to choose not to intertemporally substitute across time. When TFP is high, consumption rises and the real interest rate falls to discourage saving. If the real interest rate were constant (or if households could save) they would increase the labor supply when wages are high and smooth consumption using savings. As discussed in Chapter 13, RBC models typically assume balanced growth-consistent preferences, which imply constant labor supply along the growth path. Higher TFP has both income and substitution effects, but the assumption of linear production implies that they both cancel out in equilibrium. Constant labor supply is thus a consequence of balanced-growth preferences and consumption and output that are perfectly correlated.

Keep these results in mind when reading Chapter 18 on the New Keynesian model. The production side of that economy is identical to the one presented here. However, the introduction of nominal rigidities gives monetary policy the ability to control the real interest rate, which drives fluctuations in consumption. Wages adjust to clear the labor market. Unlike in the RBC model, where TFP shocks are the primary driver, in the New Keynesian model fluctuations in the real interest rate become the key source of output variation, operating through the intertemporal substitution of consumption.

To sum up the simple model, introducing TFP shocks and elastic labor supply is not sufficient for matching the facts on business cycles. As we will see next, with Cobb-Douglas production and capital that depreciates at a rate consistent with the data, we break the constant labor supply assumption, and generate volatile hours and investment over the cycle.

14.4.2 The core RBC model

We now move to describe the more general version of the model, referred to as the core RBC model. We will first discuss its main features and strengths, then discuss extensions developed to address some of its shortcomings. Given the thorough treatment of the neoclassical growth model in Chapter 7, our exposition will be rather brief. In particular, we focus on the planning problem, relying on the equilibrium characterization presented in Section 7.5 of that chapter. For simplicity, we abstract from population and productivity growth; alternatively, we could explicitly include growth and then transform the model into a stationary one.[10]

Preferences The economy is populated by a measure one identical households endowed with 1 unit of time that they can allocate to work or leisure. They have preferences over consumption c and leisure $1-\ell$ represented by the period utility function $U(c, 1-\ell)$. We assume that households maximize expected discounted utility, and that their preferences thus can be written as

$$\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t U(c_t, 1-\ell_t).$$

Technology Output is produced according to a production function F that takes capital k_t and labor ℓ_t as inputs:

$$Y_t = z_t F(k_t, \ell_t),$$

where z_t is stochastic TFP. F is assumed to have the usual properties, i.e., it has constant returns to scale, it is concave in both arguments, and satisfies the Inada conditions. We assume that log productivity follows an AR(1) process, $\log(z_t) = \rho_z \log(z_{t-1}) + \sigma_z \epsilon_t$, where $\epsilon_t \sim N(0,1)$ are iid random variables. Capital depreciates at constant rate δ and it takes

[10] As discussed earlier, balanced growth preferences are homothetic, so renormalizing the model under constant growth rates is straightforward.

"time to build." That is, investment in period t becomes productive in period $t+1$. The evolution of the capital stock can be written

$$k_{t+1} = (1-\delta)k_t + i_t.$$

Virtually all of the analysis for the case of the neoclassical growth model goes through nearly unchanged for the RBC model. In particular, one could show that the competitive equilibrium is efficient (i.e., the first welfare theorem holds) and that we can characterize the competitive equilibrium allocations using the solution to the social planner problem. The main difference is that allocations now depend on the history of realizations of the productivity shock. Thus, we proceed directly to analyzing the social planner problem in recursive form.

The planning problem for this economy can be written recursively as:

$$V(k,z) = \max_{k',\ell} U(zF(k,\ell) + (1-\delta)k - k', 1-\ell) + \mathbb{E}\left[V(k',z')|z\right],$$

where we've substituted out for consumption using the budget constraint. The solution to this problem delivers decision rules for capital accumulation, $k' = g(k,z)$, and one for hours worked, $\ell = h(k,z)$.

The first-order conditions for the maximization problem (assuming differentiability of the value function) are given by

$$zF_2(k,\ell) = \frac{U_2(c,1-\ell)}{U_1(c,1-\ell)} \tag{14.8}$$

and

$$U_1(c,1-\ell) = \beta\mathbb{E}[V_1(k',z')|z], \tag{14.9}$$

where U_1 is the marginal utility of consumption, U_2 is the marginal utility of leisure, and V_1 is the first derivative of the value function with respect to capital (its first argument). The envelope condition reads as

$$V_1(k,z) = \Big(zF_1(k,\ell) + 1 - \delta\Big)U_1(c,1-\ell). \tag{14.10}$$

Using this in equation (14.9) we obtain

$$U_1(c,1-\ell) = \beta\mathbb{E}\left[\left(z'F_1(k',\ell') + 1 - \delta\right)U_1(c',1-\ell')|z\right]. \tag{14.11}$$

The key optimality conditions of the real business cycle model are (14.8) and (14.11). The intratemporal optimality condition (equation (14.8)) equates the marginal product of labor (what would be the wage in the competitive equilibrium) with the marginal rate of substitution between consumption and leisure. Temporary positive shocks to z lead to increases in labor supply, as households intertemporally substitute leisure away from a time when the marginal product of working hours is high. The strength of this response depends on the Frisch elasticity, discussed in Chapter 12.

Equation (14.11) is the standard intertemporal Euler equation, which equates the marginal utility of consumption today to the expected marginal utility of consumption tomorrow, adjusted by the time discount factor β and the stochastic rate of return on capital, $z'F_1(k',\ell')+$

$1-\delta$, which in turn equals the gross real interest rate in the competitive equilibrium. To the extent that shocks to productivity are persistent, a positive shock to productivity today raises productivity tomorrow. That, *ceteris paribus*, in turn increases the marginal product of capital, which induces households to reduce consumption and increase investment tomorrow. The strength of this response depends on the intertemporal elasticity of substitution discussed in Chapters and 4 and 12.

Despite the RBC model being a fundamentally non-linear one, the original work by Kydland and Prescott (and much of the RBC literature that followed), solved the model by linearizing it around the non-stochastic steady state.[11] The linearization method is discussed in detail in Chapter 10. As noted in Section 10.6 of that chapter, linearizing the model yields linear decision rules for the aggregate variables as functions of the two states: capital k, and TFP, z, for example $y = G(k, z)$ and $k' = H(k, z)$. Using these functions, we can construct impulse response functions for the model by drawing a shock $z_0 = z_s s + \epsilon_0$ at the steady state, and then progressively applying the policy function (assuming no further shocks). This generates a sequence $y_0 = G(k_s s, z_0)$, $y_1 = G(H(k_s s, z_0), z_1)$, $y_1 = G(H(H(k_s s, z_0), z_1), z_2)$, and so on, where we've substituted the dynamics for k using H and $z_t = \rho z_{t-1}$. The sequence $\{y_0/\epsilon_0, y_1/\epsilon_0, \cdots\}$ would represent the model IRF to a one-time TFP shock. The linearized functions are essentially equivalent to a VAR representation: since the composition of linear functions is linear, we can represent a variable at time t as a linear function of lags of the variable. The model-implied impulse response can then be compared to its empirical counterpart—estimated using a VAR or local projection—to assess how well the model captures the propagation dynamics observed in the data.

IRF estimation with MIT-shocks

Analogous to how LPs and VARs estimate the same IRFs, Boppart, Krusell, and Mitman (2018) recently have shown how models can be linearized by computing the perfect-foresight impulse response of the economy to an unexpected shock—a so-called "MIT shock."[a] Instead of having a recursive linear representation of the aggregate variables as in the example above, the model is linearized in *sequence space*, where the only state variable is the time since the shock. The model can then be simulated simply by drawing shocks and superimposing them.[b] In the case of the representative-agent model, this is not necessarily advantageous numerically, since the recursive formulation only has two state variables, whereas for the impulse responses we need to keep track of t values for the impulse response. But for heterogeneous-agent models, discussed in Chapter 21, the recursive formulation has an infinite-dimensional state variable (the wealth distribution), and so linearizing the model in the sequence space a la Boppart et al. (2018) is more efficient computationally and conceptually easier to implement.

[a]Not so named because the author of this chapter is the MIT-man, but coined by Sargent (a Harvard alumnus) to criticize that type of analysis that was currently being carried out at MIT as being inconsistent with rational expectations. But, as Boppart et al. (2018) showed, if the model is linear, then solving for the "MIT shock" is fully consistent with the rational expectations equilibrium.

[11]Subsequent work showed that the RBC model is well approximated by linear approximations because the variance of TFP shocks is small.

[b] Auclert and Mitman (2018) show how to estimate and simulate models using simulated method of moments in the sequence space.

14.4.3 The Kydland-Prescott blueprint for business-cycle analysis

In addition to providing a framework for business-cycle analysis, Kydland and Prescott (1982) essentially created a blueprint for how modern, quantitative macroeconomics could be conducted. The starting point for any paper is a precise economic question. Questions can either be positive (e.g., "Can technology shocks account for the co-movement in output, consumption and investment in the data?") or normative (e.g., "How should we optimally provide unemployment insurance over the business-cycle?") in nature. Importantly, questions are about measurement and answers are numbers. The key ingredient to answering the question is a structural theory of the economy (a model). The model is a measurement device used to derive the quantitative implications of the theory. The model should be chosen from a set of "well-tested" theories, i.e., those that are based on numerous micro-econometric studies. The model can be extended to include the necessary ingredients or frictions to answer the question. Next, the parameters of the model need to be chosen, or the model needs to be "calibrated." The original blueprint of Kydland and Prescott (1982) was to calibrate the model along some dimensions of the data and then used to explain other dimensions of the data. They calibrated the model to be consistent with "long-run" facts (e.g., growth) and then evaluated it based on how well it explained business cycle facts. Finally, the model is solved numerically on the computer to solve for the equilibrium process and run the computation experiment that answers the original economic question. For example, the answer to Kydland and Prescott's original question was that shocks to TFP could explain roughly 70% of the fluctuations in output over the business cycle.

Of the steps outlined above, all should be familiar, except perhaps the "calibration" step, which we revisit here in the context of the RBC model. Calibration was first introduced Chapter 4 and then in more detail in Chapter 8, where in both cases it was applied to a simpler, deterministic environment. In this chapter, we apply the same conceptual framework to a richer setting. The general strategy is to use empirical moments at the macro level, alongside microeconomic evidence from households or firms, to discipline the model's preference and technology parameters. In Prescott's view, the purpose of calibration is not to maximize statistical fit, but rather to choose parameters that align with established economic theory and match key stylized facts. The aim is to construct a model consistent with some dimensions of the data and assess how well it explains others. Calibration also made it feasible to work with models that were, at the time, too complex for structural estimation, enabling their use in quantitative analysis.

Calibration and measurement The basic question of Kydland and Prescott (1982) was how much fluctuations in de-trended macro aggregates could be explained by "technology shocks." The idea behind the shock was one to aggregate TFP, as described in the RBC theory above. The question did not preclude other shocks from being important for business cycle fluctuations (like government-spending shocks, or shocks to people's preferences), but focused on analyzing the role of one particular shock. The question arises, however, how does one discipline the shock process to be fed into the model?

The starting point for measuring TFP was Solow's growth accounting approach, discussed in several different contexts earlier in the book (e.g., in Chapters 2 and 13), thus using the assumption of a CRS production function, perfect competition and good measurement of inputs and output, and profit maximization. Measurement of the capital stock is not immediate: the national income and product accounts typically only include measures of investment and depreciation, but not actual measures of the capital stock. The stock of capital is usually inferred from a "perpetual inventory" method. The basic idea is to make an initial guess for the capital stock, then based on measures of investment and estimates of the depreciation rate of capital δ, compute forward the time series of capital using the accumulation equation from the growth model: $k_t = (1-\delta)k_{t-1}+i_t$. A typical value estimated for δ at a quarterly frequency is 0.025.

Direct use of Solow's method involves using time-series data on capital and labor shares. Since these had not varied greatly over the postwar period (up until the mid-1980s at least), a Solow residual series based on computing the average share and then applying this share period by period did not generate large errors. We set $\alpha = 0.36$ so that it approximates the value of the capital share. Then, the TFP series z_t can be measured as

$$\log z_t = \log y_t - \alpha \log k_t - (1-\alpha)\log \ell_t.$$

Fluctuations in output at business cycle frequencies can arise from changes in inputs or from changes in TFP. By computing the Solow residual, we can gauge the extent to which technology shocks drive business cycle fluctuations. We can, in particular, estimate a stochastic process for those shocks to productivity to use as an input to the RBC model. A result is that the process for z_t is well approximated by an AR(1) process. Typical estimates on post-war data find that the persistence is quite high ($\rho_z = 0.95$ at a quarterly frequency) with a standard deviation of innovations of $\sigma_z = 0.007$.

Limitations of the approach

One of the main criticisms of this approach to calibration is that technology shocks are not measured directly; instead, they could simply be errors in the measurements of inputs and outputs (see the discussion of hard-to-observe utilization rates below). Negative shocks are challenging to understand too. On the other hand, our technologies do advance over time and the notion that these advances are not perfectly even—occur at faster and slower rates over time—is quite natural. A period of slow technological development then implies negative shocks but only relative to trend.

A related point is that multi-sectoral versions of the RBC model must have technology shocks that are highly positively correlated across sectors, since sectoral outputs (and employment) have this feature. Is such an assumption reasonable: why would the construction sector and the service sector experience positive technology shocks at the same time? Clearly, there are general technological developments—such as IT—that benefit all sectors, but it is far from clear that these dominate.

Returning to the possibility of mismeasurement, consider variable capacity utilization as well as quality improvements in the capital stock. During booms, firms would tend to use their capital more intensively, leading to higher measured TFP. Conversely, during recessions, the utilization of inputs drops, leading to lower TFP. To correct for this we can adjust capital K_t for its utilization rate u_t: $K_t^* = u_t K_t$. The composition of labor and capital can change over time. For instance, an experienced workforce or newer machinery can lead to higher output even with the same quantity of inputs. Thus, methods have been developed to adjust labor ℓ_t for changes in workforce composition, education, and experience, and to adjust capital K_t for changes in the quality and type of capital goods.

Before seeing the model in action in the next section, it's important to note that the research program undertaken by Kydland and Prescott was not designed or intended to explain business cycle fluctuations solely with technology shocks. The question was simply: how much of output fluctuations could technology shocks account for? Their original agenda included as a next step to introduce monetary features into the model and evaluate how much they contribute to fluctuations. The success of technology shocks in explaining the lion's share of business-cycle fluctuations is what ultimately led the researchers to instead increase the complexity and richness of the real setup. Given our previous discussions of the changing cyclicality of productivity post-1985, had Kydland and Prescott developed business-cycle theory in the 2000s it's quite possible that they would have continued to explore the importance of monetary (and other) phenomena.

14.4.4 The RBC model in action

After having estimated a process for TFP and specified the production technology, the final step of the calibration is to specify and parameterize household preferences. Here, again, Prescott wanted to restrict the degrees of freedom to be consistent with economic theory and empirical evidence. Given that the RBC framework was built on the neo-classical growth model, the preferences should be consistent with balanced growth. That restricted preferences over consumption to be of the CRRA form. Later, King et al. (1988) proved that preferences over consumption and leisure that were consistent with balanced growth and constant hours had to be of the form

$$\frac{\left(c_t v(1-\ell_t)\right)^{1-\sigma}-1}{1-\sigma}.$$

Here, we take v to be a strictly increasing, strictly concave function, satisfying $v'(0) = \infty$.[12] In the original calibration, Prescott argued that balanced growth with increasing wages and consumption implied a unitary elasticity of substitution between consumption and leisure, and thus specified the aggregator between consumption and leisure as Cobb-Douglas, $c^{1-\phi}(1-\ell_t)^{\phi}$. The question is then how to parameterize the shares, ϕ, the CRRA parameter, σ, and the discount factor, β. Again he turned to the implications of the growth model. In steady state the real interest rate is equal to the growth rate of consumption times σ divided by β. Then given data on the average real interest rate and consumption

[12] Boppart and Krusell (2020) later generalized these preferences to accommodate balanced growth and falling hours, consistent with recent macro data across the OECD, or hours in the U.S. from a longer-run perspective.

growth, β could be backed out as a function of σ. The value for σ was chosen based on other estimates that tried to measure it in the data based on life-cycle consumption patterns and the response of consumption and stock portfolios. These data suggested that $\sigma = 1$ was a reasonable value, implying log preferences. This then implies a quarterly value for $\beta = 0.99$, representing an annual 4% discount rate. The period utility function was then given by

$$u(c_t, 1 - \ell_t) = (1 - \phi) \log c_t + \phi \log(1 - \ell_t).$$

The final parameter to be calibrated is ϕ. One approach is to calibrate it so that the level of hours relative to the total time endowment, 1, is reasonable. There are 24 hours in the day and people typically work 8 hours, so we can then select the value of ϕ to generate $\ell_t = 1/3$. Using the first-order conditions of the model equation (14.8) and multiplying both sides by ℓ and dividing by c, with our functional forms we arrive at

$$\frac{\phi}{1-\phi} \frac{\ell_t}{1-\ell_t} = (1 - \alpha) \frac{y_t}{c_t}.$$

With α specified, and the average of y_t/c_t measurable in the data, given a target of $\ell_t = 1/3$ we can then solve for $\phi = 0.6325$.

Now that we've calibrated the model, we can solve for the dynamic system and simulate the model time series to see how they compare to U.S. business cycle statistics. Section 10.6 in Chapter 10 describes the methodology used to solve the model using linearization techniques, with Example 5 detailing a specific application to the calibrated RBC model used in this chapter. The second moments are collected in Table 14.4. All variables are treated as in the data, as long deviations from an HP filtered trend with smoothing parameter 1,600.

Table 14.4: Business cycle moments in core RBC model

Variable x	Standard Deviation	Relative Std. Dev. σ_y	Auto-Correlation	Cross correlation of x with y
Output y	1.33	1	0.73	1
Consumption c	0.42	0.32	0.81	0.90
Investment i	4.13	3.10	0.72	0.99
Hours n	0.64	0.48	0.72	0.98
TFP z	0.92	0.69	0.43	0.99

The key observation from the table is that TFP shocks successfully generate business cycle fluctuations: output fluctuates significantly in response to technology shocks. The volatility of output is larger than that of TFP, implying that the model generates amplification. Consumption, investment, and hours are all pro-cyclical, as in the data. Investment is more volatile than output, while consumption is less volatile than output, both relationships consistent with the empirical findings.

While the model is successful in generating fluctuations in output, the volatility in the model is lower than that in the data (about 1.33% in the model, compared to 2% in the 1949-1984 time period). Hours in the model are less volatile than output, whereas in the data they exhibit approximately the same volatility. This suggests that the Frisch elasticity of 2 is actually *too low* to match the fluctuations in hours worked over the cycle (which in turn would help increase volatility in output).[13] Finally, the correlation of all variables with output is higher than in the data. This is likely due to the fact that we are only considering a single impulse to the model (the TFP shock) which generates the procyclical co-movement between all variables. In reality, the economy likely experiences many kinds of shocks (we will discuss government spending shocks later in this chapter) that exhibit different co-movement patterns. Simulating the model with additional shocks could bring down the strong correlations.

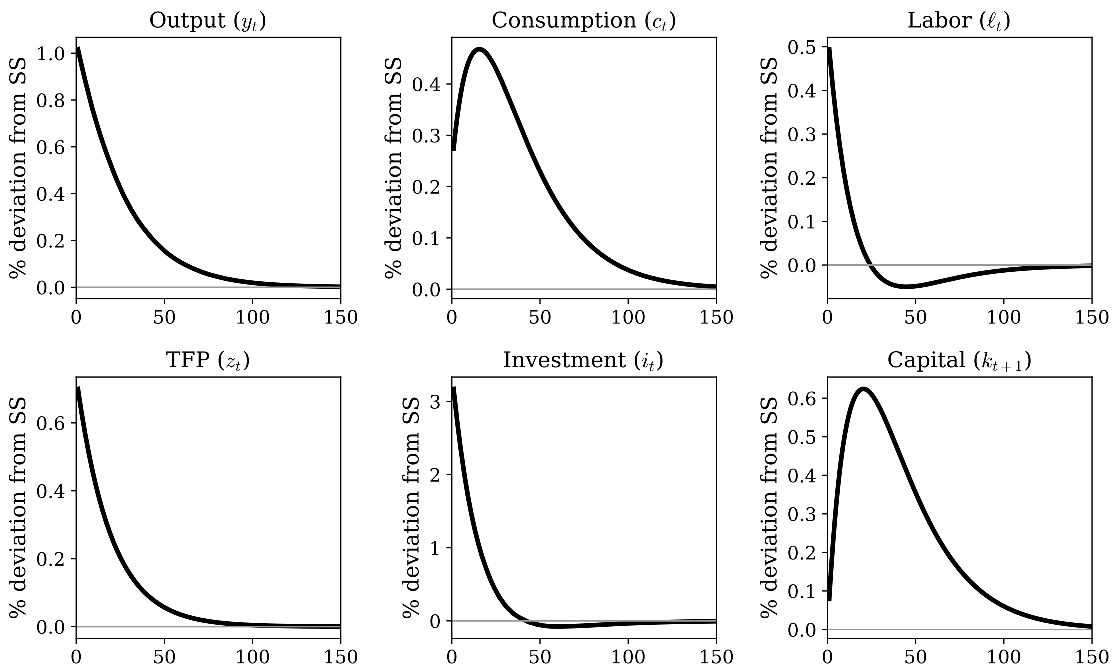

Figure 14.8: TFP shock in RBC model

The results from the core RBC model can be understood by analyzing the two key forces in the model. One is how willing households are to work more when the returns to working—measured by the marginal product of labor (MPL)—rise. The second force involves the willingness of households to smooth consumption over time in the face of fluctuating income. When a positive productivity shock occurs, the incentives to work increase (MPL rises) and output increases (about 1%). In the simple model without capital, consumption rose one-for-one with output, leaving labor supply unchanged. With capital accumulation, however, households can smooth consumption by adjusting savings. As a result, consumption increases by less than output (about 0.3%), while investment responds strongly (over 3%). Because consumption rises only modestly, the response of labor supply is more closely tied to the Frisch elasticity (here, hours increase by about 0.5% because consumption does increase). One way to think about the RBC model is that labor demand is fluctuating (as MPL moves) and moves along the labor supply curve, which in this case is roughly a line with slope equal to the Frisch elasticity. The adjustment is possible because capital accumulation allows

[13]Recall from the above discussion of the static model that the preferences used and $\ell = 1/3$ imply a Frisch elasticity of 2.

households to smooth consumption. This highlights the importance of capital for the RBC model to generate amplification and business cycle co-movement as in the data.

To better explain the mechanism in the model in response to a TFP shock, we can plot the impulse response functions in Figure 14.8. The IRFs show us in response to the impulse (the TFP shock) what happens to all of the endogenous variables in the model. By comparing how much output increases relative to TFP, we can get a measure of the amount of amplification in the model. Similarly, we can measure the propagation of shocks by looking at how long the increase in output lasts relative to the persistence of the exogenous process of TFP. Finally, the IRFs show us the shape of the responses, where they look more linear or hump-shaped. The basic RBC model generates amplification, but does not generate much actual propagation of TFP shocks. We can illustrate this more concretely by solving the model with less persistent processes for TFP. Figure 14.9 plots the impulse responses for the baseline calibration ($\rho = 0.95$) and also for iid TFP shocks ($\rho = 0$) and TFP shocks with a half-life of one quarter ($\rho = 0.5$). The persistence of output is essentially the same as the TFP process across the three impulse responses. The less persistent processes generate more amplification on impact (with hours and output responding by more), because households are more willing to intertemporally substitute their labor supply. Households respond to the temporary rise in wages by increasing labor supply, but rather than immediately consuming the additional income, they primarily save it through higher investment. This allows them to smooth consumption over time.

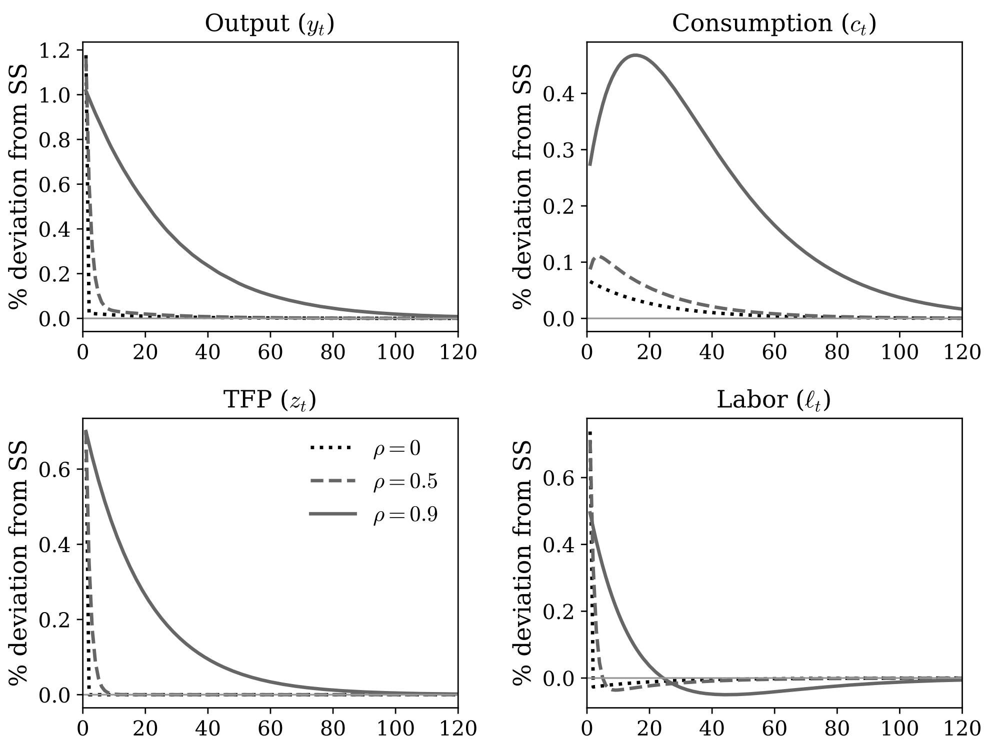

Figure 14.9: Sensitivity to TFP shock persistence in the RBC model

To further illustrate the sensitivity of the model to the two key forces, we explore sensitivity to the intertemporal elasticity of substitution (recall with CRRA preferences the EIS is $1/\sigma$) and to the Frisch elasticity. In the baseline model we had an EIS of 1 and a Frisch elasticity of 2. We can recalibrate the model with higher and lower Frisch elasticities (0.5 and 10), and higher and lower EIS (0.5 and 2), and compare the IRFs to those in the baseline model. In Figure 14.10 we plot the IRFs for those alternate calibrations.

Focusing first on the left panel, if we increase the Frisch elasticity, households are more willing to intertemporally substitute labor supply, and we get more amplification. Hours increase by more in response to the TFP shock, which leads to higher output, and more volatile investment as households save the extra earnings to smooth consumption going forward. With a lower Frisch elasticity we get the opposite, more dampening. The core RBC model gave a response of hours and output that was too low relative to the data, suggesting that a calibration with a higher Frisch elasticity may improve the fit. However, at the same time that leads to more volatile investment, which was already too high relative to the data.

Turnings to the right panel of Figure 14.10, a higher intertemporal elasticity of substitution also leads to more amplification of output and labor supply in response to TFP shocks. On impact, however, the response of consumption, is non-monotonic in the EIS (it actually falls in both cases). When households value consumption smoothing less, they are more willing to intertemporally substitute to take advantage of the productivity shock. Households are willing to increase consumption by less today and lower leisure more to increase investment to take advantage of the persistently higher TFP. That leads to a bigger boom in consumption in the medium run, after househoulds have built up the capital stock. Then, as TFP returns towards steady state, households begin eating down the extra capital that they had accumulated. When households are less willing to intermporally substitute consumption, they substitute more on the leisure margin, as can be seen by the solid yellow line in the figure. Hours rise on impact, but quickly fall below steady state levels.

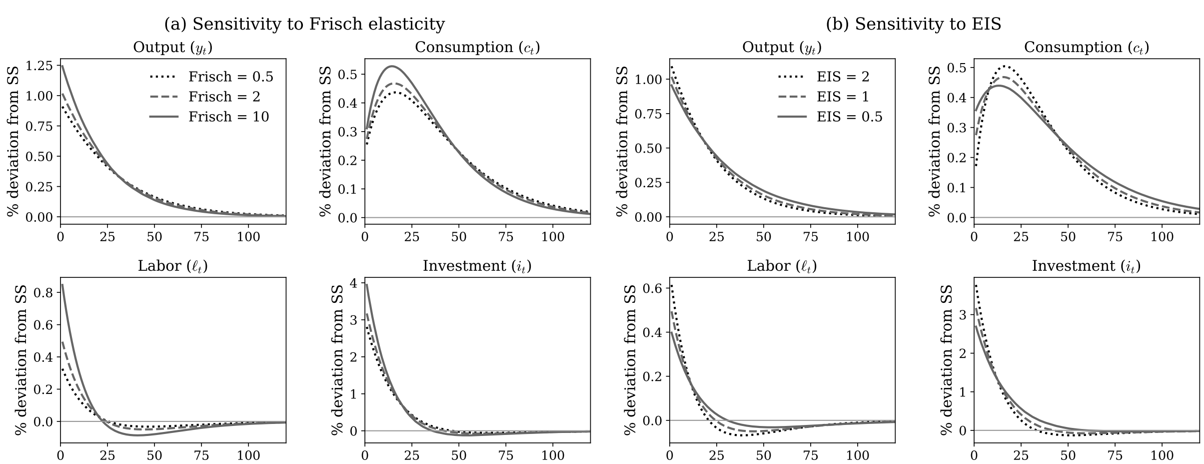

Figure 14.10: Sensitivity to EIS and Frisch Elasticity in RBC Model

To sum up, evaluated in the context in which the original RBC papers were published, the model was very successful at generating business cycle fluctuations and co-movements. The model explained nearly 80% of fluctuations in output, and replicated the relative volatilities of investment and consumption. That being said, the model still featured too much consumption smoothing and hours that were not volatile enough. Several extensions to the core framework (some of which are discussed later in this chapter) were developed to address some of the shortcomings of the model.

14.5 Extensions to the RBC Model

In this section, we briefly discuss some of the most prominent extensions to the RBC model, many of which remain as the building blocks for modern business cycle analysis.

14.5.1 Indivisible labor

One of the earliest extensions of the RBC model introduced indivisible labor, addressing a key critique: the assumption that households can adjust work hours smoothly at business-cycle frequencies. This frictionless adjustment was increasingly questioned by applied microeconomists. As discussed earlier in this chapter the implied Frisch elasticity by the calibration was approximately 2. Early empirical estimates using micro data (e.g., MaCurdy, 1981 and Altonji, 1986a) suggested the Frisch elasticity was close to zero.[14] The "micro" Frisch elasticity appeared small, but at the macro level the Frisch elasticity needed to be much higher to match the data.

A second critique was that there was no notion of unemployment or inactivity: everyone was always fully employed at the hours that she wished to work. At the same time, the data suggest that much of the aggregate hours adjustment in the data actually occurs along the extensive margin—because of fluctuations in aggregate employment—not because of hours conditional on being employed. Hansen (1985) and Rogerson (1988) introduced indivisibilities into the labor supply choice. They made the simple assumption that a consumer could either work full time, supplying $\bar{\ell}$ hours, or not work at all. The discrete choice introduces potential non-convexities, which would complicate the dynamic programming problem of the household. To simplify the analysis, they assumed that people had access to lotteries to convexify their labor supply choice. As discussed in Chapter 12, the model essentially collapses to one with an infinite aggregate Frisch elasticity at the extensive margin, since the disutility of hours is linear. The model also features a low elasticity of labor supply at the intensive margin—zero by construction in this case—which helps reconcile large fluctuations in employment with relatively small changes in hours worked among those employed.

14.5.2 Capital adjustment costs

A second shortcoming of the core RBC model was that investment was too volatile relative to output. One way to dampen the investment response to TFP shocks is to introduce capital adjustment costs to explain gradual investment behavior (Hayashi, 1982). Costs represent the real resource costs incurred by firms when changing their capital stock and are meant to capture various factors, including installation and setup costs, costs related to training workers to use new technologies, and costs associated with selling old equipment.

[14] One limitations of these early studies was that they focused on prime-age males. Women and younger and older men tend to have more elastic labor supply.

The investment adjustment cost shows up in the capital accumulation equation,

$$k_{t+1} = (1-\delta)k_t + i_t - \Psi(i_t, k_t),$$

where $\Psi(i_t, k_t)$ is a function capturing capital adjustment costs. Typically, we assume that adjustment costs are convex (to keep the optimization problem well behaved), and take a quadratic form:

$$\Psi(i_t, k_t) = \frac{\psi}{2}\left(\frac{i_t}{k_t} - \delta\right)^2 k_t.$$

The idea is that maintain the capital stock at its previous level by replacing the depreciated capital is costless (e.g., maintenance as opposed to buying new machines).[15] The quadratic formulation yields a shadow price of "installed" capital equal to

$$1 + \psi\left(\frac{i_t}{k_t} - \delta\right),$$

which we define as q_t: Tobin's q. Firms increase investment when the price of capital exceeds that of output (or consumption), and reduce it otherwise.

When firms face these costs in adjusting their capital stock, they will spread their investment decisions over time. This can lead to more persistent and amplified responses to shocks, helping the model better match the observed persistence and magnitude of business cycles. New Keynesian models with capital also tend to feature adjustment costs to dampen the response of investment to movements in the real rate.

14.5.3 Variable capacity utilization

Variable capacity utilization was introduced by Greenwood et al. (1988) and further explored by others. The basic idea is that firms can choose how intensively to use their capital stock depending on the current prices for their goods. You can imagine that machines could be run for multiple shifts of workers, up to 24 hours a day. Or they could be used for one eight hour shift and then sit idle the other sixteen hours of the day. Variable capacity utilization allows for more flexible responses to shocks, as firms can adjust capital usage without immediate investment. The effective capital in the production function is thus modified as $u_t k_t$, at the cost of higher depreciation for higher capacity utilization:

$$k_{t+1} = (1 - \delta(u_t))k_t + i_t.$$

The idea is that if you utilize your capital more intensively today, then it will depreciate faster. Variable capacity utilization allows another margin of adjustment in response to TFP shocks. Now, firms can increase their effective capital services in addition to hiring more labor. This leads to greater amplification of TFP shocks, as compared to the standard model. The measurement of TFP shocks, however, now needs to be adjusted, since our Solow residual approach assumed a constant depreciation rate in constructing the capital stock series.

[15] Note also that the formulation is homogeneous of degree 1 in (i, k), so as to be consistent with zero profits for firms.

14.5.4 Additional shocks

In the standard RBC model, there is only one shock, z_t. As a result, all variables are highly correlated with each other. One way to break the high correlation between all of the variables is to introduce additional shocks. Here we discuss two of the most common additional shocks in the RBC literature, investment-specific technology shocks and government spending shocks.

Investment-specific technology shocks These shocks affect the efficiency of new capital goods. They separate the technology affecting consumption goods production from the technology affecting investment goods production, allowing for differential productivity growth rates. The concept was introduced in a series of papers by Greenwood et al. (1997); Greenwood, Hercowitz, and Krusell (2000) who looked at both the long-run and business cycle implications of investment-specific change; the long-run part is discussed in Chapter 13. Proposing investment-specific productivity shocks was motivated by the negative correlation between the price and quantity of investment in equipment at long-run and business cycle frequencies (the opposite of what one would expect from the adjustment cost model discussed above). In addition, investment-specific shocks can break the strong correlation between consumption and investment in the core model. The capital accumulation equation changes as follows:

$$k_{t+1} = (1-\delta)k_t + q_t i_t$$

, where q_t is the investment-specific shock. They found that investment-specific shocks can account for roughly 30% of business cycle fluctuations.

Government spending shocks Another common shock introduced is a government spending shock. The shock can help break the high correlation between wages and hours, in which is at odds with the data. Let g_t indicate per capita government expenditure, so we modify the resource constraint from the core model to read:

$$c_t + i_t + g_t = y_t.$$

Why does introducing government consumption help break the correlation between hours and wages? Suppose households care about an aggregate consumption $C_t = [\theta g_t^{\varphi} + (1-\theta)c^{\varphi}]^{1/\varphi}$, where φ measures the elasticity of substitution between c and g. If the elasticity is very high, government spending is like private consumption, and it will not have any effect on the outcomes we've looked at so far. On the other hand, if this elasticity is low, government spending is more like a tax and it reduce available resources, increasing labor supply due to a negative income effect.

14.6 Business cycle accounting

In the previous section, we explored several frictions and extensions to the RBC model aimed at improving its empirical performance. While each brought the model closer to the data, it remains useful to have a systematic framework for assessing which features of the model align with the data and which do not. In this section, we introduce such an approach. Chari, Kehoe, and McGrattan (2007) propose Business Cycle Accounting (BCA) as a methodology that can be used to assess how different economic frictions (such as technology shocks, labor market frictions, or fiscal policy variations) contribute to business cycle dynamics. We can think of the frictions as deviations from the behavior that would be implied by the core RBC model. Their motivation was to provide a structured way to evaluate different frameworks (including extensions to RBC) based on their ability to explain observed economic fluctuations. The approach introduces four broad potential frictions into the RBC model: efficiency wedges (technology), labor wedges (labor market), investment wedges (capital investment), and government spending wedges (fiscal policy). These wedges can be interpreted as distortions away from the planner's allocation (recall that the equilibrium in the core RBC model is efficient) that could arise from various frictions or policies.

The wedges amount to modifications of the equations in the dynamic programming problem faced by the households that show up as potential distortions:

1. **The efficiency wedge (z_t)**: This wedge affects the production function, which shows up as a TFP shock:
$$Y_t = z_t F(k_t, \ell_t).$$

2. **The labor wedge ($1-\tau_t^\ell$)**: This wedge shows up in the household's budget constraint as a tax on labor earnings and distorts the labor-leisure choice and can be seen as distortions in the labor market. Equation (14.8) is modified to read
$$(1-\tau_t^\ell) z_t F_2(k_t, \ell_t) = \frac{U_2(c_t, 1-\ell_t)}{U_1(c_t, 1-\ell_t)}.$$

3. **The investment wedge ($1/(1+\tau_t^i)$)**: This wedge, which multiplies investment in the budget constraint, affects capital accumulation equation and can be interpreted as frictions in the investment sector. Equation (14.11) is thus specified as
$$U_1(c_t, 1-\ell_t)\left[1+\tau_t^i\right] = \beta \mathbb{E}\left[\left(z_{t+1}F_1(k_{t+1}, \ell_{t+1}) + (1-\delta)\left[1+\tau_{t+1}^i\right]\right) U_1(c_{t+1}, 1-\ell_{t+1}) | z_t\right].$$

4. **The government consumption wedge (g_t)**: This variable affects the resource constraint and can be seen as reducing output (and in this sense is a distortion), given that utility is not affected by it.
$$c_t + i_t + g_t = Y_t.$$

In addition, we have to modify the household budget constraint with lump-sum transfers T_t, so that the aggregate resource constraint holds in equilibrium.

In practical terms, the efficiency wedge can capture a variety of phenomena, including technological changes, market power, changes in regulatory environments, and other factors that affect how efficiently an economy can convert inputs into outputs. In the framework, it will show up identically to a change in TFP that we've considered thus far. The labor wedge captures discrepancies between the marginal rate of substitution between consumption and leisure and the marginal product of labor. Factors contributing to the labor wedge include search frictions, taxation (see Chapter 15), and other policies affecting labor supply and demand, such as unemployment insurance or minimum wage laws. A significant labor wedge indicates that the labor market is not operating efficiently, potentially leading to underemployment or an inefficient allocation of labor resources. The frictions underlying the labor wedge will be explored further in Chapter 20, which introduces frictional labor markets and search.

The investment wedge is a measure of any distortions affecting the relationship between the cost of investment and the return on investment. It can arise from financial frictions (see Chapter 19), taxation on capital income (see Chapter 15), and policies affecting savings and investment decisions. An adverse investment wedge means that investment in the economy is lower than what would be expected based on the fundamentals of the economy, leading to underinvestment in productive capital and, consequently, slower economic growth.

Finally, the government spending wedge affects the allocation of resources between private and public sectors. It reflects the impact of government expenditure (and taxation policies to fund this expenditure) on the economy's productive efficiency and the private sector's consumption and investment decisions. A government spending wedge might arise from government spending either crowding out private investment by raising interest rates or by reallocating resources in ways that are not aligned with market efficiency.

In order to implement the BCA approach, we first calibrate the preference and technology parameters of the model. Here, Chari et al. (2007) follow a similar procedure to the one we discussed above for calibrating the core RBC model, using the same functional forms and choices for the intertemporal elasticity of substitution and Frisch elasticity. Next, they assume that the four aggregate shocks (the wedges) follow an AR(1) process. Then they log-linearize the model with respect to the aggregate state and estimate the stochastic process for the wedges via maximum likelihood comparing the model generating time series for output, hours, investment, and government consumption from both the data and the simulated time series generated by the model. Armed with the estimated time series and stochastic process for the wedges, they can then evaluate the consequences of the various wedges in isolation (or combinations of the four).

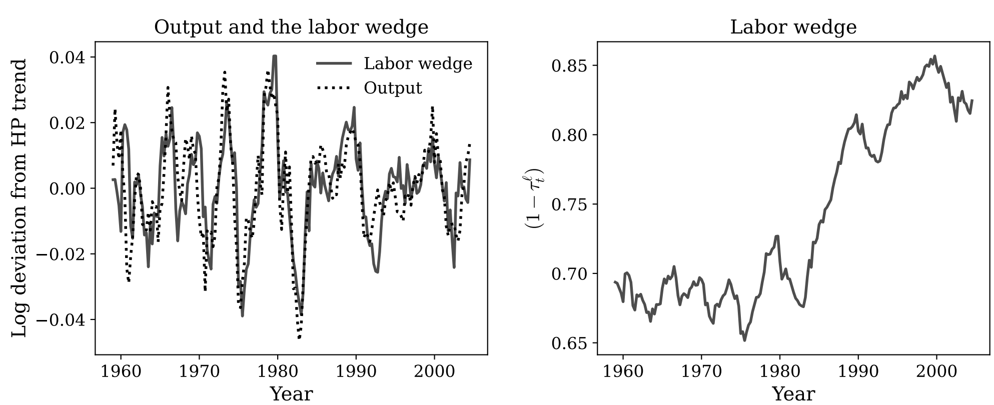

Figure 14.11: The estimated labor wedge from BCA

The outcome of the BCA estimation is that the efficiency wedge and the labor wedge are the two most important wedges for driving business cycle fluctuations: about 70% of the variance in output fluctuations are driven by the efficiency wedge and 25% by the labor wedge. These findings hold even when the authors look at the Great Depression (the main exercise was based on post-war data). The findings suggest that the potentially most promising extensions to the RBC model are ones that introduce explicit frictions in the labor market. The estimated labor wedge in log-deviations from an HP-filtered trend and in levels are plotted in Figure 14.11. Focusing first on the left panel, we can see that the labor wedge and output are highly correlated and have nearly the same standard deviation over the estimation sample: when output is low, the "tax" on working is high, making the wedge high: it is as if people like working less in recessions. Clearly, labor-market frictions and involuntary unemployment is a plausible interpretation of this finding, but not the only one. Turning to the right panel, the level of the labor wedge has increased significantly over time (meaning less distortions in the labor market). The timing of the secular change in the labor wedge coincides with the changing nature of fluctuations that we saw in the data post-1985.

Recent work on "jobless recoveries"

The labor wedge itself explains nearly 75% of the variance in fluctuations in labor, suggesting that the labor wedge is potentially related to the phenomena of "jobless recoveries" that we discussed earlier in the chapter. What could explain both the cyclical and secular trends in the labor wedge? Two recent theories have emerged that explain "jobless recoveries" based on real phenomena. One, proposed by Mitman and Rabinovich (2019), posits that countercyclical and unemployment-dependent extensions of unemployment benefits act as a time-varying policy wedge. Using a framework with frictional labor markets similar to that developed in Chapter 20, they show that benefit extensions explain roughly 1/3 of post-war labor market dynamics, and can explain the emergence of jobless recoveries (even though the underlying parameters of the model and shock processes remain constant).

Another explanation points to a striking trend in the U.S. around the same time period: namely the increase in female labor-force participation. The increase was particularly pronounced for married women. A series of recent papers (Olsson et al., 2019; Fukui, Nakamura, and Steinsson, 2023; Albanesi, 2019) thus show that for men, even the recessions before 1991 were "jobless", but masked in the aggregate by increasing female labor-force participation along the extension margin. The secular increase in participation can be explained by the decline in the gender-wage gap, which would mimic the secular trend in the estimated labor-wedge. When the secular increase in female participation plateaus in 1990, the jobless recoveries emerge in the aggregate. These facts can be rationalized in an extension to the RBC model with indivisible that takes gender and household composition seriously.

These examples help illustrate why the core of business cycle research has moved beyond the standard RBC model. Frictions in the labor market and household heterogeneity have emerged as important features in explaining business cycle phenomena (see, e.g., Krusell et al., 2017, Krueger, Mitman, and Perri, 2016). These will be discussed in Chapters 20 and 21, respectively.

14.6.1 Current frontiers of business cycle research

More than 40 years after the publication of Kydland-Prescott, business-cycle research continues to thrive. The COVID-19 pandemic, ensuing recession, and subsequent global surge in inflation have made clear that the period of tranquility leading up to the Great Recession—known as the Great Moderation—is over. As predicted by Lucas (1980):

> "One would expect developments to arise from two quite different kinds of forces... Of these forces the most important... consists of purely technical developments that enlarge our abilities to construct analogue economies... The second source of technical developments is changes in the questions we want models to answer..."

Researchers have continued to extend the core RBC framework to meet the economic challenges of the 21st century. In the decade since the Great Recession, we have seen progress in incorporating household heterogeneity into workhorse macro models. One of the critiques of modern macro was its focus on the representative agent. While real business cycle theory was built on micro foundations, the frameworks were inconsistent with extensive empirical micro evidence on household behavior. The research frontier is now focused on making the core RBC theory more *micro-consistent* as a framework, i.e., consistent with empirical microeconomic evidence on household and firm behavior, expectations, and outcomes (for consumption, an idea pioneered by Deaton 1992). It combines incomplete markets at the household level, and frictional product and labor markets.